An Interop Publication

ConneXions®

The Interoperability Report

July 1993 Volume 7, No. 7

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Mountain View, CA 94040–1219, USA. +1 415-941-3399. Fax: +1 415-949-1779. Toll-free: 1-800-INTEROP. E-mail: connexions@interop.com.

ISSN 0894–5926

From the Editor

The *Simple Network Management Protocol* (SNMP) has been the subject of several articles in *ConneXions.* An important aspect of the SNMP framework is the concept of a *Management Information Base,* or MIB. A MIB contains definitions of the managed objects for a particular device, expressed in a subset of the ASN.1 data representation language.

A *MIB compiler* is a tool used by authors of SNMP MIBs as well as implementors of SNMP agents and users or SNMP applications. In addition to verifying the syntactic correctness of a MIB, a MIB compiler can automatically generate data structures and code required by an agent to implement a particular MIB. Our first article this month is by Sam Roberts of Farallon Computing and provides an introduction to SNMP MIB compilers.

Three months ago I attended the IPNETWORKING '93 conference and exhibition in Birmingham, England. The event reminded me in many ways of an early INTEROP—in the days when there were only a few dozen vendors on hand, and all conference sessions ran in a single track. There is something about this kind of informal workshop atmosphere that fosters cooperation and a genuine willingness by all parties to perform interoperability testing, and that is exactly what happened in Birmingham. Six router vendors demonstrated PPP and OSPF across the exhibition network. Ron Catterall and Jacqui Holland-Bradley, the organizers of IPNETWORKING '93, describe the tests in an article starting on page 10.

Perhaps the most exciting aspect of networking today is in the area of applications development. Particularly interesting are the number of distributed information tools such as *WAIS, WWW, Archie, Gopher,* and *Netfind,* as described in previous issues of *ConneXions.* The common denominator for all these systems is that they provide a *service* to end *users* of the network. Last November, a conference was held in Pisa, Italy, to bring together developers and users of these and other network tools. Jack Kessler attended this *Network Services Conference* and brings us a report on page 16.

From our somewhat expanded announcements section this month, I'd like to draw your attention to a couple of items: First, both the President and [a subset of] the House of Representatives are now reachable via electronic mail. Second, make sure you grab a copy of the latest *Network Reading List* developed by Charles Spurgeon at the University of Texas at Austin.

Finally, next month is INTEROP, so call us today at 1-800-INTEROP or 1-415-941-3399 to register or request more information. There are still a few available *Birds Of a Feather* (BOF) slots for INTEROP. If you'd like to organize a BOF, send me e-mail: ole@interop.com.

An Introduction to SNMP MIB Compilers

by Samuel M. Roberts, Farallon Computing, Inc.

[Ed. The following article is reprinted with permission from *The Simple Times,* Volume 2, Number 1, 1993.]

Introduction

An SNMP *MIB compiler* is an extremely useful tool, not only for authors of SNMP MIBs, but also for implementors of SNMP agents and users of SNMP management applications. In addition to verifying the syntactic correctness of a MIB, a MIB compiler can automatically generate data structures and code required by an agent to implement a particular MIB. A MIB compiler can also make information about managed objects in proprietary vendor MIBs or new Internet-standard MIBs available to a management application.

This article provides an introduction to SNMP MIB compilers. It begins with a description of the function that a MIB compiler performs. It then examines the internal structure of a typical MIB compiler and provides an overview of some of the more common output formats produced by SNMP MIB compilers. Finally, the article concludes with a section describing a number of openly-available and commercial MIB compilers.

What is a MIB compiler?

A compiler is a program that translates a program written in one language—the *source* language—into an equivalent program in another language—the *target* language. Typically, the source language is a high-level programming language, such as C or *Pascal,* while the target language is a low-level programming language, such as a particular target platform's assembly or machine language.

The source language for a MIB compiler is the language *Abstract Syntax Notation One* (ASN.1). ASN.1 is not a programming language; it is a language for describing structured information. ASN.1 resembles the data declaration portion of a high-level programming language. The input to a MIB compiler isn't a program, but rather a collection of MIB modules written in a subset of ASN.1. These MIB modules contain definitions of managed objects that correspond to information in network devices that can be manipulated via SNMP.

MIB compilers usually generate various alternate representations of the managed object definitions in the input MIBs. These alternate representations are easier for management applications and agents to process than the ASN.1 representation.

Some alternate representations are actually data structure declarations in a high-level programming language, such as C, that can be compiled and linked into a management application or agent. Others are data files containing "flat" representations of the managed object definitions that can be read into memory by a management application or agent at run-time.

In some cases, MIB compilers output code to assist in the implementation of the input MIBs. For example, a MIB compiler may generate skeleton routines for retrieving or setting the value of a managed object, or routines to generate particular SNMP Trap-PDUs.

Structure of a MIB compiler

Despite the fact that a MIB compiler accepts input written in a data description language rather than a programming language, it is similar in structure to a programming language compiler. It typically consists of a *front-end* and one or more *back-ends.*

The front-end reads input files containing ASCII text corresponding to ASN.1 MIB modules and constructs an internal representation of the information contained in the MIB modules. The back-end produces one or more output files from this internal representation.

The front- and back-ends of the MIB compiler may be totally independent applications that communicate by means of an intermediate file on disk, or they may be individual modules of a single application that communicate by means of shared data structures in memory.

Although most MIB compilers are implemented as stand-alone programs and generate output files on disk, a MIB compiler may sometimes be integrated directly into a network management application (or even an agent). In this case, the MIB compiler consists only of a front-end and no back-end, and the MIB compiler does not produce an output file. Instead, the management application makes direct use of the internal representation constructed by the front-end from the ASN.1 MIB modules.

The Front-end

The front-end of a typical MIB compiler consists of a "lexical" analyzer, a "syntactic" analyzer, and a "semantic" analyzer, together with support routines such as symbol table management and error handling routines. The lexical analyzer, or scanner, reads the sequence of ASCII characters in the input files and groups them together into a sequence of *tokens*. Tokens are the basic lexical units of ASN.1. There are keyword tokens, such as "BEGIN" and "END," punctuation tokens such as "::=" and ".." identifier tokens such as "iso," and numeric tokens such as "255."

The *syntax* analyzer, or *parser*, takes the tokens produced by the scanner and groups them together into syntactic structures according to the ASN.1 grammar. The token stream produced by the scanner must consist entirely of one or more instances of ASN.1's `ModuleDefinition` syntactic structure; otherwise the input contains syntax errors and the parser generates appropriate error messages. For example, the parser would recognize the following sequence of tokens:

```
INTEGER {
        up(1),
        down(2),
        testing(3)
}
```

as a syntactically correct instance of the `IntegerType` syntactic structure using the following *Backus-Naur Form* (BNF) productions (or rules) from the ASN.1 grammar:

```
IntegerType ::=
      INTEGER
    | INTEGER "{" NamedList "}"

NamedList ::=
      Named
    | NamedList "," Named

Named ::=
     identifier "(" Signed ")"

Signed ::=
     number
   | "-" number
```

continued on next page

Introduction to SNMP MIB Compilers *(continued)*

Parse tree

To see that the previous sequence of tokens is indeed an instance of the `IntegerType` syntactic structure, it is helpful to construct a diagram called a *parse tree*. A parse tree exhibits the syntactic structure of a particular sequence of tokens.

The parse tree for the previous example is as follows:

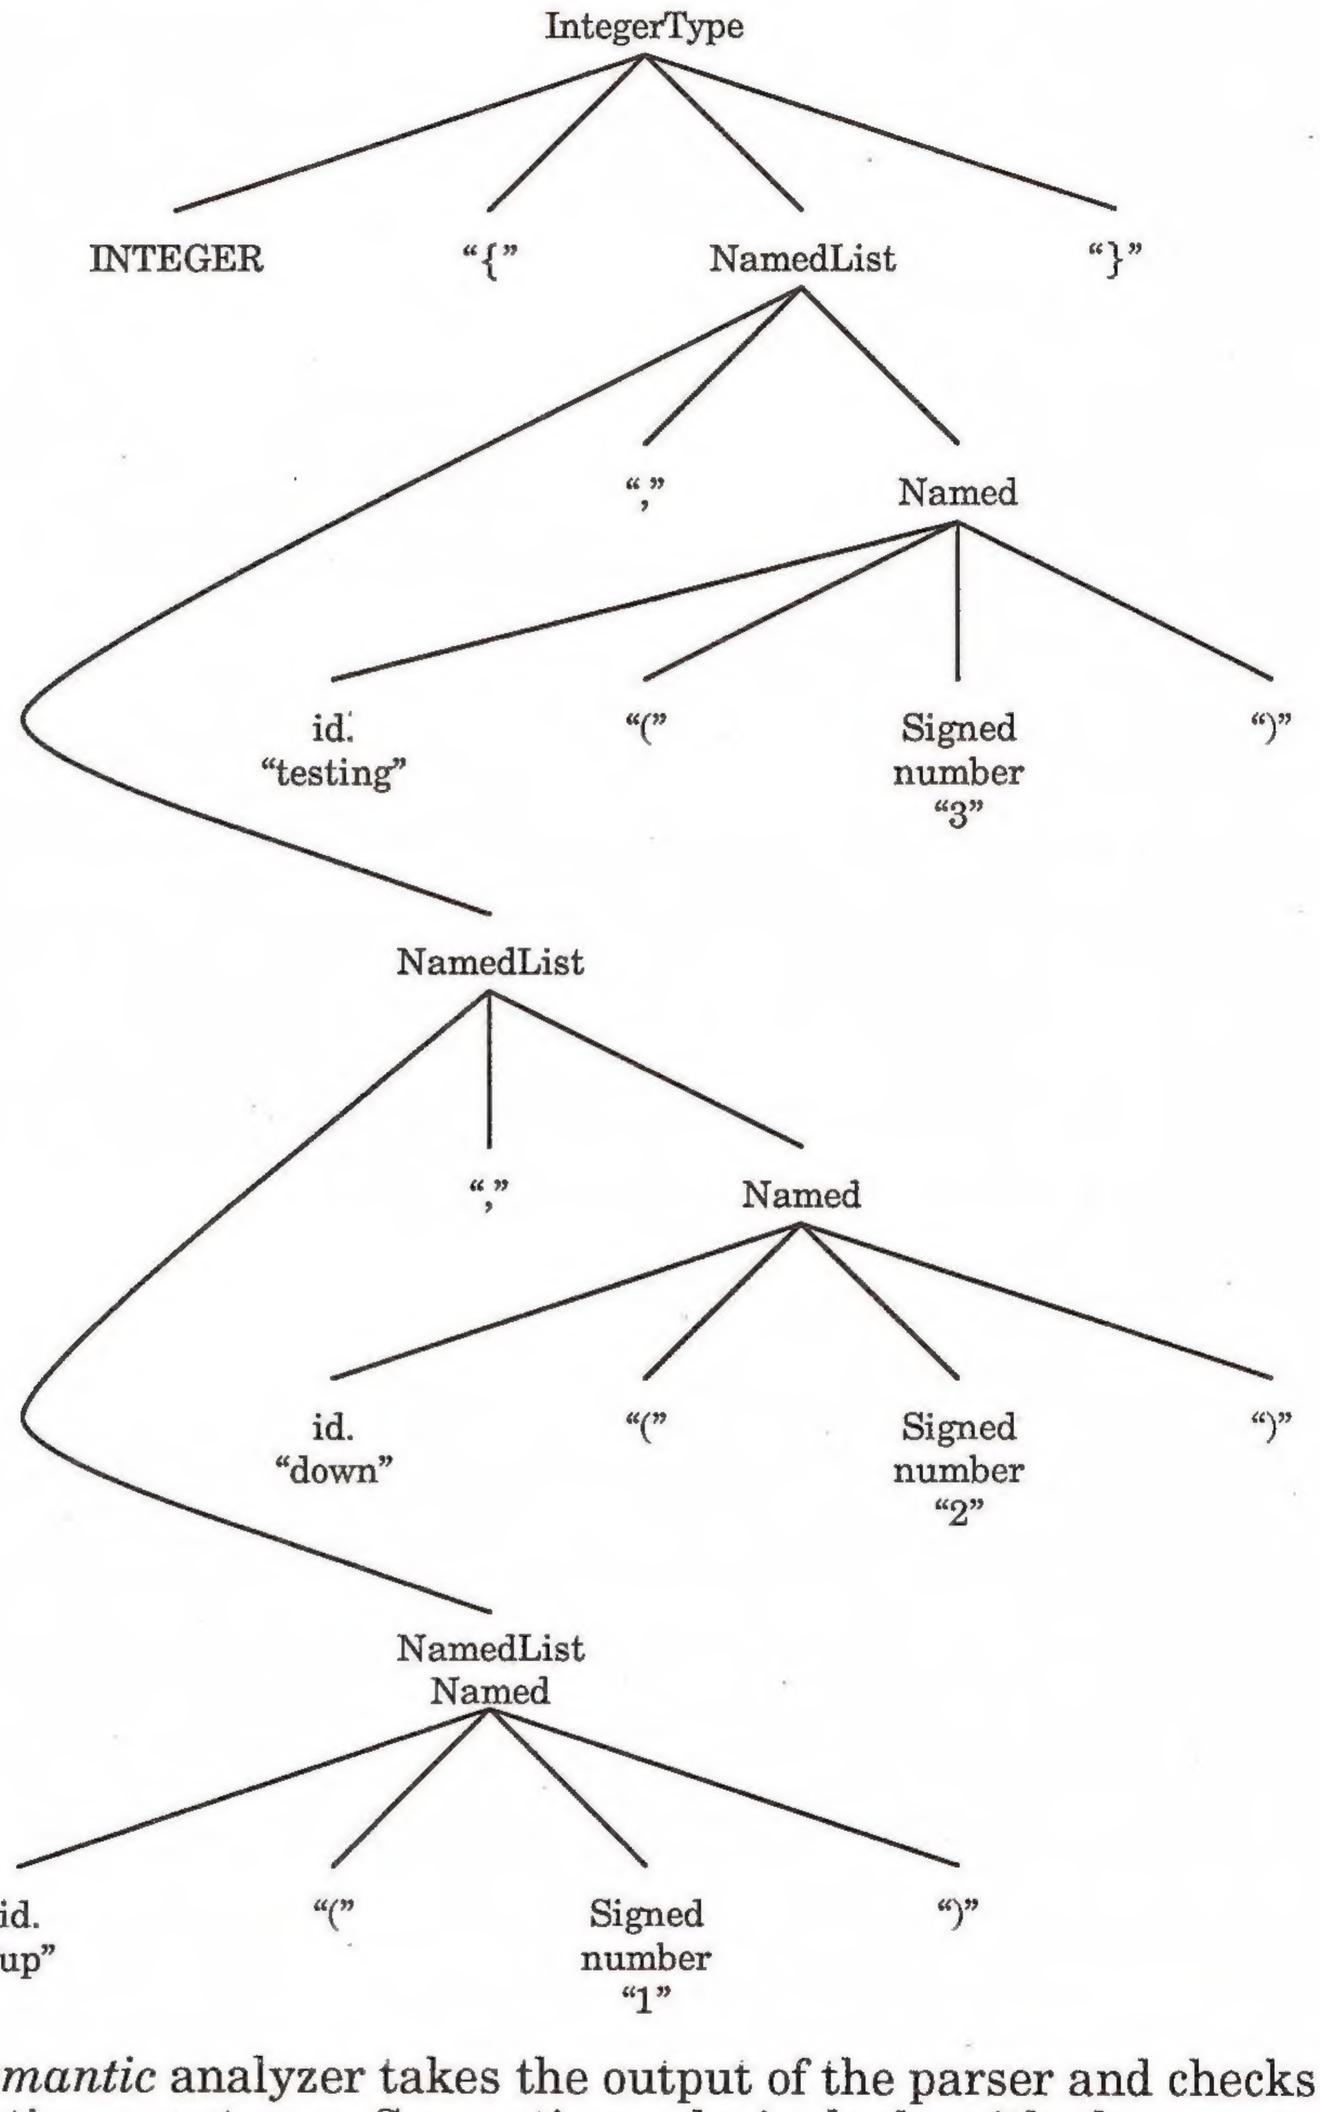

The *semantic* analyzer takes the output of the parser and checks it for semantic correctness. Semantic analysis deals with those aspects of ASN.1 that cannot be specified by means of a grammar, but instead depend on context. Even though a particular sequence of tokens may be syntactically correct according to the ASN.1 grammar, it may not be legal ASN.1. For example, the sequence of tokens:

```
MyType ::= INTEGER
```

is always a syntactically correct instance of a `TypeAssignment`. However, depending on the context in which this `TypeAssignment` appears, it may not be semantically correct. For example, it would be semantically incorrect if the enclosing MIB module contained a prior assignment to the type identifier `MyType`.

Error handling

The lexical, syntactic, and semantic analyzers each make use of various support routines, including error handling and symbol table management routines. Error handling routines generate error messages for the user when there is an error in the input to the compiler. Errors can be generated at any phase of the compilation. For example, the scanner will generate an error if it encounters a character that cannot be part of any ASN.1 token, such as an at-sign, or an improperly formed token, such as an identifier that ends with a hyphen. The syntax analyzer will generate an error if it encounters a sequence of tokens that isn't valid according to the ASN.1 grammar, such as a comma in the middle of an OBJECT IDENTIFIER value. The semantic analyzer will generate an error if it detects a semantic error, such as a conflict between the syntax specified for an object in a SEQUENCE definition and a later OBJECT-TYPE definition.

Symbol table management routines keep track of the names used in the MIB modules being compiled and record essential information about each name, such as its type (textual convention, managed object, trap, and so on) and the modules in which the name is defined.

Note that although conceptually it is useful to think about these front-end components as independent tasks, they do not necessarily operate sequentially. For example, the same routine that parses a particular syntactic construct often checks it for semantic correctness. In addition, the scanner is often implemented as a subroutine of the parser, which asks the scanner for the next token when the parser needs one.

Back-ends

Many MIB compilers contain multiple back-ends, each of which produces an output file containing a different representation of the same set of input MIBs. If the front- and back-ends of the MIB compiler are integrated in a single application, then one typically uses command line switches to select a particular output format. If the front- and back-ends are implemented as separate applications that communicate by means of an intermediate file, then one generates the desired output format by executing the appropriate back-end after generating the intermediate file using the front-end.

The output from some back-ends is specific to the implementation of agents, while the output from others is primarily suitable for use by management applications.

Back-ends that aid in agent implementation

Before describing the output from back-ends designed to aid in the implementation of SNMP agents, it is useful to have a basic understanding of the structure of a typical SNMP agent. Most SNMP agent implementations separate SNMP protocol processing from MIB variable access. All knowledge of the SNMP protocol, including the formats of the various *Protocol Data Units* (PDUs), and the encoding and decoding of ASN.1 data, resides in an SNMP protocol engine. All knowledge of the managed objects in the MIB resides in a set of access functions commonly known as "method routines."

Typically, a **C** data structure representation of the MIB tree provides the glue between the protocol engine and the method routines. When the protocol engine receives an incoming SNMP request, it extracts the OBJECT IDENTIFIER name of the object instance whose value is to be retrieved or changed. Using the OBJECT IDENTIFIER prefix that forms the name of the corresponding OBJECT-TYPE, the protocol engine consults the **C** data structure to find pointers to the method routines specific to that OBJECT-TYPE. The protocol engine invokes the appropriate method routine to perform the requested operation on the specified object instance.

continued on next page

Introduction to SNMP MIB Compilers *(continued)*

Constructing the **C** data structure representation of a MIB tree by hand, even once, is tedious at best. Several MIB compilers provide back-ends that generate the required **C** data structure automatically from a set of input MIBs. (Some agent implementations do not link the MIB tree into the agent at build time. Instead, these implementations generate the tree structure at run-time from a data file containing a "flat" representation of the MIB tree. A MIB compiler is normally still used to generate this data file, albeit using a different back-end.)

Several MIB compilers contain back-ends that assist in the implementation of the method routines for an agent. One of these back-ends generates a **C** source file containing skeleton function definitions for each method routine. These skeleton functions are typically stub functions with the appropriate return type and argument list. The agent writer fills in the code required to get or set the appropriate MIB variable. A clever compiler may even be able to fill in some or all of the code required to implement a particular method routine.

Another back-end produces a **C** header (`.h`) file containing ANSI **C** and non-ANSI **C** function prototypes for the method routines. This header file is included by the file containing the MIB tree data structure and the file containing the definitions of the method routines.

At least one MIB compiler contains a back-end that produces a **C** source file containing functions that an agent can invoke to send SNMP Trap-PDUs corresponding to the `TRAP-TYPE` definitions contained in the input MIBs. These functions include code to create variable bindings in the Trap-PDU for the variables specified in the `VARIABLES` clause of the `TRAP-TYPE` definition.

Back-ends for use with management applications

Most MIB compilers include one or more back-ends that generate output files designed for use by management applications. One output file format common to a number of MIB compilers is known as *mosy format* after the name of a popular MIB compiler, *mosy*. A *mosy format* file includes a flat representation of the MIB tree that management applications can use to map between descriptors and the corresponding OBJECT IDENTIFIERS. It also specifies each object's syntax, access, and status. The following example is an excerpt from a *mosy format* file generated from RFC 1213 (MIB-II):

```
mib-2          mgmt.1
system         mib-2.1
...
sysDescr     system.1 OctetStringread-only mandatory
sysObjectIDsystem.2 ObjectID   read-only mandatory
...
```

Several compilers include back-ends that output extended versions of *mosy format*. These formats preserve more of the information contained in the input MIBs, such as indexing information for objects in conceptual rows and subtype information for objects with range and size constraints. They also include additional information for objects with enumerated integer syntax to enable management stations to map integer values to the corresponding identifiers.

MIB compiler availability

This section describes those MIB compilers which are the most popular and with which the author is most familiar. Mention of a particular compiler does not imply its endorsement, either by the author or by *The Simple Times*. (Or by *ConneXions*).

There are at least two openly-available stand-alone MIB compilers. The first is named *mosy,* which is an acronym for "Managed Object Syntax-compiler (YACC-based)." This compiler is a part of the *ISO Development Environment* (ISODE), a package for developing OSI protocols and applications. It is available by anonymous FTP from `ftp.uu.net` in the directory `/networking/osi/isode`.

The second, named *SMIC* (the "SNMP MIB Compiler"), generates output files in *mosy* and extended *mosy* format. It also produces a more complex *SMIC*-specific format meant for use with as yet unreleased back-ends. *SMIC* is available by anonymous FTP from `ftp.synoptics.com` in the directory `/eng/mibcompiler`. (Volume 1, No. 4 of *The Simple Times* contains an announcement for *SMIC.*)

In addition, Carnegie Mellon University (CMU) distributes an openly-available SNMP implementation. The CMU SNMP package does not contain a stand-alone MIB compiler. Instead, it provides a library that management applications can call at startup to compile MIB modules from disk into memory. Although the CMU SNMP package can also be used to implement an SNMP agent in a managed device, no stand-alone MIB compiler is available for generating the **C** data structure representation of the agent's MIB tree. The CMU SNMP package is available via anonymous FTP from Internet host `lancaster.andrew.cmu.edu` in the directory `/pub/smp-dist`.

Both Epilogue Technology, Inc. and SNMP Research, Inc., market commercial SNMP implementations in source code format that developers can license to incorporate SNMP functionality into management stations or network devices. Each company licenses a MIB compiler designed to work with its SNMP implementation.

The Epilogue Technology MIB compiler includes a variety of back-ends that produce MIB representations suitable for both agents and management stations. Other back-ends generate **C** source code that can assist in the implementation of method routines and the generation of SNMP Trap-PDUs.

The SNMP Research MIB compiler consists of two independent programs. The first is an enhanced version of the openly-available *mosy* MIB compiler described above. Its back-end produces an enhanced *mosy format* output file. The second program processes the enhanced *mosy* output and generates a variety of output files suitable for implementing both management stations and agents.

In addition, many commercial network management applications include a MIB compiler for importing information about managed objects, particularly those in vendor MIBs. For example, SunConnect's UNIX-based network management station, *SunNet Manager,* includes a stand-alone MIB compiler named *MIB2Schema* that translates MIBs into a proprietary format data file (called a *schema file*) that is read by the *SunNet Manager* application. Similarly, Hewlett-Packard's *OpenView Network Node Manager* application contains an integrated MIB compiler that can be used to import information about managed objects directly from MIB modules.

Choosing a MIB compiler

The intended use of a MIB compiler may affect the criteria used in choosing one. For example, if the goal is to compile MIBs for use with an agent or management application, then it's desirable to use a MIB compiler that can successfully parse MIBs containing some minor ASN.1 syntax errors. Otherwise, a fair bit of time may be spent correcting syntax errors in MIBs, since even some MIBs published as RFCs contain syntax errors.

continued on next page

Introduction to SNMP MIB Compilers *(continued)*

The *mosy* MIB compiler, and some compilers derived from *mosy,* perform less strict syntactic and semantic checking than other compilers, and as a result they will successfully compile MIBs with some syntax errors. Other compilers perform stricter syntactic and semantic checking, and provide command line switches that invoke less stringent checking.

On the other hand, if you are writing a proprietary vendor MIB, or a new Internet-standard MIB, you probably want a MIB compiler with very strict checking that will catch any errors you inadvertently introduce. *SMIC* and the Epilogue Technology MIB compiler are particularly good for this purpose.

Finally, most of the MIB compilers mentioned in the previous section are available in source code format, with the exception of those compilers that come with commercial network management applications. As a result, many of the compilers can be modified to generate output in formats they don't already support. Although a compiler designed to generate output for one SNMP implementation can be modified to generate output files for a different one, it is usually best to use the MIB compiler that was designed for use with a particular SNMP implementation or management application.

References

[1] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Introduction to version 2 of the Internet-standard Network Management Framework," RFC 1441, April 1993.

[2] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1442, April 1993.

[3] Galvin, J. M., & McCloghrie, K., "Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1445, April 1993.

[4] McCloghrie, K. & Galvin, J. M., "Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1447, April 1993.

[5] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1448, April 1993.

[6] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1450, April 1993.

[7] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Manager-to-Manager Management Information Base," RFC 1451, April 1993.

[8] *ConneXions,* Two *Special Issues* on Network Management and Network Security, Vol. 3, No. 3, March 1989 and Vol. 4, No. 8, August 1990.

[9] Case, J., McCloghrie, K., Rose, M. T., Waldbusser, S., "The Simple Management Protocol and Framework: Managing the Evolution of SNMP," *ConneXions,* Volume 6, No. 10, October 1992.

[10] Marshall T. Rose, *The Simple Book—An Introduction to Management of TCP/IP-based internets,* Prentice-Hall, 1990, ISBN 0-13-812611-9. (Second edition due out in August 1993).

SAM ROBERTS holds an A.B. from Dartmouth College and a J.D. from the University of California at Berkeley. Since 1986 he has worked on a variety of networking-related projects for Farallon Computing, a leading supplier of networking software and systems for Windows and Macintosh personal computers. He is a co-editor of the 802.3 MAU MIB. He can be reached as:
`sroberts@farallon.com`

Information on *The Simple Times*

Contents

The Simple Times (ISSN 1060-6068) is an openly-available publication devoted to the promotion of the *Simple Network Management Protocol.* In each issue, *The Simple Times* presents: a refereed technical article, an industry comment, and several featured columns:

- Applications and Directions
- Security and Protocols
- Working Group Synopses
- Ask Dr. SNMP
- Standards

In addition, some issues include brief announcements, summaries of recent publications, and an activities calendar. Past technical articles have included:

- A New View on Bulk Retrieval with SNMP
- Sets are Fun: Introducing the SMDS Subscription MIB Module
- An Implementation of SNMP Security
- Customer Net Management of the InterSpan Frame Relay Service
- Accomplishing Performance Management with SNMP

Policy

The Simple Times is openly-available. You are free to copy, distribute, or cite its contents. However, any use must credit both the contributor and *The Simple Times.* Further, this publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *The Simple Times.*

Subscription information

The Simple Times is available via electronic mail in three editions: *PostScript,* MIME (the multi-media 822 mail format), and richtext (a simple page description language). For more information, send a message to: `st-subscriptions@simple-times.org` with the word "help" in the Subject: field. In addition, *The Simple Times* has numerous hard-copy distribution outlets. Contact your favorite SNMP vendor and see if they carry it. If not, contact the publisher and ask for a list:

The Simple Times
c/o Dover Beach Consulting, Inc.
420 Whisman Court
Mountain View, CA 94043-2186
Tel: +1 415-968-1052 • Fax: +1 415-968-2510
E-mail: st-editorial@simple-times.org

Submission information

The Simple Times solicits high-quality articles of technology and comment. Technical articles are refereed to ensure that the content is marketing-free. By definition, commentaries reflect opinion and, as such, are reviewed only to the extent required to ensure commonly-accepted publication norms. *The Simple Times* also solicits announcements of products and services, publications, and events. These contributions are reviewed only to the extent required to ensure commonly-accepted publication norms.

Submissions are accepted only in electronic form. A submission consists of ASCII text. (Technical articles are also allowed to reference encapsulated *PostScript* figures.) Submissions may be sent to the contact address above, either via electronic mail or via magnetic media (using either 8-mm *tar* tape, $^1/_4$-in tar cartridge-tape, or 3-$^1/_2$in MS-DOS floppy-diskette).

Each submission must include the author's full name, title, affiliation, postal and electronic mail addresses, telephone, and fax numbers. Note that by initiating this process, the submitting party agrees to place the contribution into the public domain.

The Interoperability Test at IPNETWORKING '93

by
Ron Catterall & Jacqui Holland-Bradley, IPNetworking, Ltd.

Introduction

Interoperability and open systems are the underlying requirements of most builders and users of computer networks. The days of being tied into a single vendor are just about over, and IT departments are expecting to exploit the commercial advantages of choice of vendor. At the networking protocol level, the *Internet Protocol Suite* (IPS), more commonly referred to as "TCP/IP," provides the "glue" to bind a multi-vendor network together. These standards, which grew out of the *Department of Defense Advanced Research Projects Agency* (DARPA) initiative in the 1960s, have become the de facto standard for communications software, and form the basic protocol supporting the world-wide Internet which currently has more than 1.7 million host computers attached.

The communications architecture binding this huge Internet together is built out of leased telecommunications lines joining switching computers commonly known as "Gateways" (original DARPA nomenclature) or "Routers" as they are now more usually known. This long distance data communications structure is composed of autonomous networks which are very often built out of routers from a single vendor: e.g., the JIPS academic network is Cisco based, ICNET is Proteon based etc. The reasons for this single vendor approach in the standards-based, open world of TCP/IP are often largely ones of convenience, but there is also an area of concern about the ability of different implementations of the IP protocol suite to interoperate over serial (telecommunications) lines. These concerns are mainly concentrated in the implementations of the *Point to Point Protocol* (PPP) used to communicate between routers over serial lines. PPP is a relatively new protocol. Another problem area is the relatively new *Open Shortest Path First* (OSPF) routing protocol used by the routers to achieve optimal transport efficiency across meshed networks.

Aims

The aim of the interoperability tests at the IPNETWORKING conferences is to enable router vendors to demonstrate the conformance of their products to these IPS standards known as PPP and OSPF, and to show the openness of their products to an audience of network builders and managers.

The old sales approach of "buy my product: it is better because it has different properties to those of the competition" is being replaced by "buy my product: it is better because it conforms to the (IPS) standards closely enough to interoperate with those of the competition." The ultimate failure from a sales point of view is now a failure to interoperate.

At the TCP level, interoperability across Local Area networks (LANs) has been almost complete for many years: it is now quite difficult to find TCP/IP implementations on host computers or on routers which do not interoperate across an Ethernet. Similarly, the older *Routing Information Protocol* (RIP) is sufficiently mature to ensure interoperability between vendors. Proprietary routing protocols are clearly beyond the pale when it comes to interoperability.

Thus the primary objective of the test is to get vendors to show compatibility of their implementations of OSPF and PPP: to do this in a realistic environment requires the use of serial lines rather than LAN links between routers.

Design

To demonstrate interoperability between routers thus requires a complex mesh of telecommunications lines: a full demonstration between the six router vendors taking part in the test requires a mesh of 15 serial lines. OSPF interoperability can then route traffic via the path of least resistance across any combination of any number of the 6 products. Equally, of course, it is not possible in practice to determine which route was taken by any particular packet in a given transmission. Nevertheless, the complete assembly is required to demonstrate the full interoperability at the OSPF level. In order to show full interoperability it is also necessary to demonstrate the passage of traffic across all possible pairs of routers individually: this guarantees the complete interoperability of the PPP implementations of all the vendors.

We define interoperability between routers A and B as the ability to pass messages from workstations attached to different LANs connected via a serial line between two routers from different vendors:

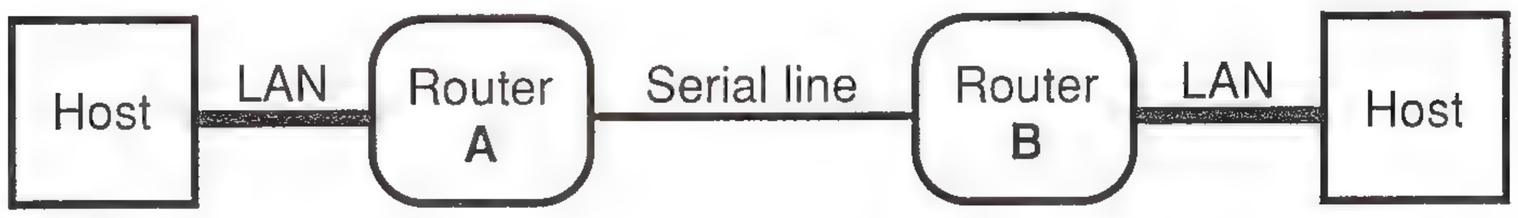

This test was carried out for all possible pairs of routers. In practice, the test involved the ability to achieve remote login between a local host (a Compaq 4/120S running the Netmanage Chameleon TCP/IP Dll under Windows 3.1) and a Sun workstation based at the Network Information Center (NIC) in the USA. Thus in practice, several other routers existed between router B and the final destination, but these remained constant throughout the tests. The link to the Internet was achieved by a bridged ISDN link between the Conference Centre and ICNET. ISDN bridges were kindly loaned at very short notice by Sonix and Alpha networks.

Since it is necessary, during the tests, to eliminate all other possible alternate routes (to prevent OSPF from finding acceptable routes around non-interoperating links), it was necessary to re-configure the topology of the mesh of serial lines for each of the 15 tests. The flexibility to do this reconfiguration was achieved using a *Bytex Unity Matrix Switch* configured to emulate 64k kilostream links between each individual pair and between all possible pairs of routers. The Bytex switch was kindly loaned to the conference by Bytex Datacom Ltd. and proved to be a very reliable and flexible piece of equipment which made the extensive re-configurations of the test network possible.

In addition to the PPP and OSPF interoperability testing, a very limited test of interoperability between a single pair of routers across a serial line using the *Intermediate System to Intermediate System Intra-Domain Routing Protocol* (IS–IS) was also carried out.

Participants

The participants in the IPS tests were: Cisco, Proteon, Timeplex and Wellfleet (full PPP and OSPF testing), NAT and Digital (PPP and RIP testing only: neither vendor supporting OSPF at the present time). In addition Digital and 3Com demonstrated IS–IS interoperability.

(At the specific request of a conference delegate, 3Com and Proteon also demonstrated PPP and OSPF interoperability, but this was not part of the planned testing.)

continued on next page

Interoperability Test at IPNETWORKING *(continued)*

Testing strategy

The interoperability tests were carried out under the control of John Hopkins of IPNetworking Ltd. with adjudication by Ole Jacobsen, Editor and Publisher of *Connexions—The Interoperability Report,* USA.

All products included in the interoperability testing were required to be standard products currently on sale at the time of the test. The specific version numbers of the software products is detailed in the table.

Results

The matrix of results as authenticated by all vendors in the presence of the adjudicator is given in the table opposite. Of the four vendors included in the full PPP and OSPF testing (Proteon, Wellfleet, Cisco and Timeplex) only one failure to interoperate was observed.

Proteon and Timeplex were unable to demonstrate complete interoperability. The problem is outlined in footnote 2 to the table on the next page, and is being followed up by both vendors, although the problem appeared to lie in the occasional faulty response of the Timeplex software to 56 byte *Internet Control Message Protocol* (ICMP) messages generated by Proteon.

NAT (who did not have a production version of OSPF available at the time of the test) demonstrated full RIP and PPP interoperability in all tests with Proteon Wellfleet, Cisco, Digital and Timeplex.

Digital (who did not have a production version of OSPF available at the time of the test) demonstrated full RIP and PPP interoperability with Proteon, Wellfleet, Cisco and NAT, but failed to achieve full interoperability of PPP and RIP with Timeplex. The problem appeared to lie in the lack of support by Timeplex for an IPCP "config.request" with "IP Address" from Digital.

Digital and 3Com achieved full IS–IS interoperability, the only vendors to attempt this test.

(Although 3Com declined to take part in the full interoperability testing, an ad-hoc, unplanned test of PPP and OSPF interoperability between Proteon and 3Com was undertaken successfully at the specific request of a conference delegate.)

About IPNetworking

History

IPNetworking was formerly the *United Kingdom Internet Consortium,* UKIC, which was formed in late 1990 to encourage the provision of public access internetworking in the UK, and onward to worldwide locations. In late 1992 UKIC felt that they could be proud that the aim of providing or seeing provided a public UK internet connected to the world wide Internet had now been realised, and felt that the structure as a Consortium was inappropriate to the continued support of the IP Users Group. IPNetworking was formed to take over the activities, (primarily the User Group and related events) in an endeavour to ensure that IP users in the UK had representation of their aims and ideals.

Activities

IPNetworking runs an *IP Users Group* (IPUG) which is the primary technical forum. The IPUG serves to inform technical and management staff of issues and developments. IPNetworking also organises technical workshops and conferences which include exhibitions and User Group meetings.

	PROTEON V 12.0C	WELLFLEET V 5.76	CISCO V 9.1	NAT V 2.10	DIGITAL V 2.0	3COM V 6.0	TIMEPLEX V 3.1.3
PROTEON		OSPF PPP	OSPF PPP	RIP PPP	RIP PPP	PPP OSPF	See note 2
WELLFLEET	OSPF PPP		OSPF PPP	RIP PPP	RIP PPP	NOT TRIED	OSPF PPP
CISCO	OSPF PPP	OSPF PPP		RIP PPP	RIP PPP	NOT TRIED	OSPF PPP
NAT	RIP PPP	RIP PPP	RIP PPP		RIP PPP	NOT TRIED	RIP PPP
DIGITAL	RIP PPP	RIP PPP	RIP PPP	RIP PPP		IS–IS and integrated IS–IS See note 3	See note 1
3COM	PPP OSPF	NOT TRIED	NOT TRIED	NOT TRIED	IS–IS and integrated IS–IS See note 3		NOT TRIED
TIMEPLEX	See note 2	OSPF PPP	OSPF PPP	RIP PPP	See note 1	NOT TRIED	

Tested at 3 pm on Thursday 1st April 1993.

Notes Accompanying Interoperability Testing

NOTE 1 = Timeplex/DEC IPCP Problem: DEC sends IPCP "config.request" with "IP Address" which Timeplex does not appear to support. We believe that an IPCP "option reject" should include the option being replied to. RFC 1332 & RFC 1331 to be consulted for verification.

NOTE 2 = Timeplex/Proteon Problem: Proteon sending ICMP echo requests to Timeplex, 56 byte length. Timeplex occasionally ignores 4 of such packets and generates 4 truncated IP datagrams in response.

NOTE 3 = Digital/3Com: Testing included RIP, IS–IS and Integrated IS–IS testing over both FDDI and PPP links. Both IP and OSI end system connectivity over the Integrated IS–IS network links were tested. As the 3Com Netbuilder was connected to the main test network through the DECNIS 600+ route propagation from RIP to IS–IS and IS–IS to RIP was enabled in order to propagate routing information from the rest of the test network to the 3Com Netbuilder.

NOT TRIED = 3Com were invited specifically to test IS–IS and Integrated IS–IS interoperability with Digital, and as such were not expected to participate in the main Interoperability area. Results of last years testing are available from 3Com.

Bytex: The Unity Matrix switch was configured to emulate 64k kilostream links between any router vendor, no problems were encountered on any of the sync links. The T160 switching hub was configured for Ethernet functionality, with router vendors switched on or off the main segment without problems. In both cases no interoperability problems were observed and the manipulation of both sync and Ethernet connections were centrally managed from the management platform.

Cabletron Systems: monitored the interoperability test with the Spectrum Network Management system. Alarms were shown while routers were being re-configured. Comprehensive management of the network was then re-established.

Dawson and Gibbons: As part of the installation, all data lines were fully tested with no faults arising. During the interoperability tests it was observations confirmed that the data lines remained stable.

continued on next page

About IPNetworking *(continued)*

Membership

Membership of IPNetworking is more likely to be appropriate to a senior member of management with an interest in public policy and the development of a national infrastructure. IPNetworking members include banks, airlines, a nationwide group of solicitors, insurance and medical companies, research organisations and universities.

Policy

The UK needs access to the same communications facilities as those in other countries, to be able to communicate worldwide with the fast-growing Internet community. It is important to redress the advantages enjoyed by competitors with a proactive approach to networking.

Services to members

IPNetworking provides the following services to its members:

- *Quarterly Journal:* IPNetworking focuses on issues and events relevant to all IP users. Besides acting as a channel for information about User Group activities, IPNetworking offers news of who is using IP and how, significant announcements from equipment vendors and service providers, and analysis of the impact of changes in internetworking technology in the UK and worldwide.

- *Bulletins:* Key developments in internetworking are covered in Bulletins to members, giving early warnings of movements in standards and trends in worldwide network operations.

- *Conference/Exhibition:* The annual conference and exhibition is the only UK event specifically for IP users and suppliers. The conference offers keynote speakers of world repute as well as practical sessions describing hands on experience of building IP networks. The exhibition is a chance to see the offerings of all major UK service and equipment providers under one roof, as well as attracting US suppliers.

- *Workshops:* As Internet access and internetworking grows in the UK there is a need to train staff not as yet familiar with the required principles and practice, and to update the knowledge of existing staff. The User Group will be running regular courses to meet these needs.

- *Access to Consultancy Services:* Substantial discounts on occasional consultancy can save you your membership fees in the first week. IPNetworking is building up a list of approved consultants, and is always interested in hearing from individuals or organisations with IP skills.

- *Help Desk:* IPNetworking provides a help desk service which endeavours to provide answers to members' questions relating to internetworking such as network contacts, applications for IP addresses, supplier contact details and electronic mail addressing. When an organisation is just starting to connect its networks together this can be a valuable and time-saving resource.

Benefits of membership

- *Advance Information on Equipment and Futures:* Organisations supporting mission-critical systems need as much advance information as possible in order to ensure their plans are in line with forthcoming developments. We seek to ensure that our members receive accurate and up-to-date information on product and service futures. You can plan your networks with the latest developments in mind.

- *Feedback from other Major Users:* It can often be helpful to know what other similar organisations are planning, or what has already been achieved in an area of technology where you are formulating your own plans.

The User Group helps and encourages its members to exchange information where this does not breach confidentiality or jeopardise competitive advantage. You benefit from awareness of general trends in usage and early reports on newly installed network services and equipment.

• *Influencing UK Internetworking Affairs:* At present internetworking in the UK is controlled by disparate organisations acting in a relatively uncoordinated and unaccountable manner. The services are sometimes available and sometimes not. The User Group seeks to ensure a coherent infrastructure for UK IP users, such as has already been developed in the US. By joining the User Group you add your weight to the "market pull" for interoperability and direction within the internetworking community, and getting the services and infrastructure you need.

Membership categories

The following membership categories exist:

• *Honorary Membership—By Board Invitation:* The purpose of honorary membership is to allow key industry individuals and organisations throughout the world to offer participation and support, thus raising the profile of the User Group in the eyes of the organisations it wishes to influence.

• *Corporate Membership:* Corporate members must be organisations who do not provide internetworking services or equipment as revenue generating activities. Their mission-critical systems will depend crucially on internetworking technologies. They are therefore often likely to be multi-site organisations, often the UK arm of multinational organisations. All services described above are available to corporate members.

• *Trade Membership:* Trade membership is open to companies whose business is the provision of internetworking services or equipment. The User Group will offer them an opportunity to discuss matters of importance to major IP users in the UK as well as providing all the services described above.

• *Individual Membership:* The Users Group offers one-day meetings on topical technical subjects, with opportunities to discuss current equipment, services and particular problems. Journal subscription is included. A series of technical seminars and workshops for members offers the opportunity to bring knowledge up to date.

For further information, please contact: IPNetworking Limited, 10 Gate Street, London WC2A 3PH, England. Ph.: +44 071 404 6771.

RON CATTERALL gained his PhD and DSc for work on localised electron states in dielectric media. Currently he is Head of the Central Computer Unit at the Imperial Cancer Research Fund. In 1989 he built ICNET, the UK's first major national network based on routed TCP/IP protocols. He was chairman of the Board of the UK Internet Consortium, and a founder member of the EBONE Consortium of Contributing Organisations. He is currently a director and Board member of IPNetworking Ltd, a UK company organizing the IPNETWORKING Conferences and Exhibitions and the IP Users Group. He provides consultancy services in IP networking to major UK corporations. He can be reached as: `ron@omega.lif.icnet.uk`

JACQUI HOLLAND-BRADLEY came to internetworking via a career in law and personnel. In 1990 she joined The UK Internet Consortium, a group of enthusiasts who wished to see the establishment of commercial Internet connectivity in the UK. In early 1992 she established a User Group for IP Users, and in October 1993 when UKIC was dissolved, she formed a new Company under the name of IPNetworking becoming the CEO. The new Company assumed all the committments of UKIC, and extended the facilities available to members to include the establishment of an annual IP conference and exhibition in the UK, workshops, help desk facilities, and publication of a quarterly journal of which she is editor. She can be reached as: `j_holland@icrf.icnet.uk`

Europe—at least—discovers the users

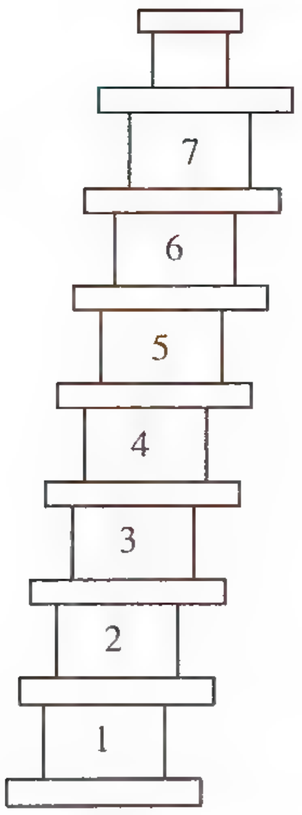

by Jack Kessler, AKCO Inc.

The "OSI Seven-Layer Model," as seen by a pretty rowdy crowd of networkers, librarians, and Italian waiters at the banquet of the Network Services Conference at Pisa, November 2–4, 1992...

...Hey, it was late at night, and we'd all worked long and hard, and the food was good, and there was lots of wine...

Conferences in Europe are a little different from conferences in other places. Ever try to build ten different six-foot towers-of-Pisa out of place cards, while "O Sole Mio" is crashing out from the tables in the back?

Introduction

The idea of a "Network 'Services' Conference" came to Europe at the beginning of last November. It hasn't arrived in the US or elsewhere, yet, and in all cases is a bit overdue. There was much of interest for anyone who loves libraries, books, networked information, or European ways of looking at things, and for anyone blessed or cursed with the need for working on a computer.

The Pisa conference brought some leading lights from North American information networking—Peter Deutsch, creator of *Archie* (not the comic book, although he carries that with him), and John "Matrix" Quarterman—together with network leaders from all over Europe, to discuss what to do about a new topic: *the users*. There were many librarians there: most of us were left fascinated, but also shaking our heads and groaning. It seems that the great amount of work so far done to help users on the networks leaves much still to be done, in both Europe and elsewhere.

Sessions

The conference, sponsored by EARN (*European Academic and Research Network*) and a group of several other organizations, attracted 360 participants, from 46 countries, and by all accounts was highly provocative and successful. Sessions covered "New Global Information Tools" (*World-Wide Web, WAIS, Gopher, Hyper-G, Archie* and the *Soft Pages Project*), "Beyond ASCII" (imaging, and ISO standards), "The Electronic Library" (projects in Israel and France, "The Virtual Library," Project *PegUn/Janus* at Columbia University), "Delivering Messages to the Desktop," "Central and Eastern Europe," "User Support," "Special Interest Communities" ("Electronic Pierce," biology, chemistry, "Human Genome"), "Managing Network Information Services," and "Information Overload." It was for me a very different European version of the birthpangs of this technology's application.

Opening Keynote

The first keynote speaker, Peter Deutsch, delivered a fascinating and funny talk—speaking at his accustomed rate, described as "56k with no flow-control"—about the necessity now for "building networks, not just network links," for "real services, not just projects," and for "not explaining, but hiding" FTP and the various other user's tools so far developed. "Nobody ever wanted a $^1/_4$ inch drill bit," he asserted, "they wanted a $^1/_4$ inch hole." The time has come, he said, to provide real information on the networks, and not just tools for getting there.

Tools

Deutsch distinguished four purposes for existing network tools:

- Class Discovery—more tools are needed, he said,
- Instance Location (indexing tools)—we have lots now,
- Access (FTP, etc.)—lots,
- Management of Information (*WAIS, 3W/WWW* at CERN)—we could use some more.

The tools and projects which exist, he said fall into four groups:

- Interactive Message Systems (*telnet, rlogin, talk, chat,* MUDD—Multi-User Dangerous Dragons),
- Store-and-Forward (e-mail, news),
- Information Delivery (Anonymous FTP, *Gopher, 3W/WWW, Prospero, WAIS, ALEX*)—the point being now to begin hiding these, hide the network, make it transparent, and,
- Tools for Finding Things—Peter's "own particular sandbox" at the moment, he says, in which he's finding that, "a gigabyte no longer is that big of a deal."

But the networks will be "useful only if populated with useful information," Deutsch said. "Librarians," moreover, "should be running the networks, not the UNIX weenies." He is concerned about the latter's penchant for reinventing the wheel first developed by the former. "It's going to be services," he concluded, "if someone around you starts talking technology, watch out."

Imaging projects

Anne Mumford, in "Beyond ASCII," pointed out that the problem with images arriving now is their use, rather than the more technical problems of their storage: image users will want to cut and paste, insert, catalog, index, and change formats, just as they now do with ASCII, she said. She mentioned CARL's Group 3 fax format journal project, "CORE—Chemistry On-line Retrieval Experiment" which stores the page and ASCII and a picture caption index, Northern Telecom's "CGM—Computer Graphics Metafile" format, and Elsevier's project for issuing 35 imaged journals on CD-ROM.

Standards: first round

Borka Jerman-Blazic described the Herculean/Augean effort currently going in to develop international standards for software. The world has over 3000 spoken languages, she pointed out, over 100 of these written: 50% use the Latin alphabet, but the other 50% use over 23 different alphabets, counting only those which have over 1 million users. So users come to the networks familiar with Latin diacritic and non-diacritic alphabets, non-Latin alphabets (Cyrillic, Greek), diacritical scripts (Arabic, Hebrew), and syllabic (Kanna Japanese) and ideographic (Chinese) written modes of expression. One might just make them all learn American English, but then again they might not want to, and they simply might not. ISO 10646, a standard on which she's working, specifies over 65,000 characters in world languages: she bravely asserted both that it will accommodate UNICODE, and that conforming commercial products will begin to appear next year.

The French libraries and ILL

Christine Deschamps delivered an elegant overview, of the vast array of current events in France. She described their work on a national ILL "union catalog": SQL request handling, an X.25 ILL system which batches requests, and a project to develop an "OSI/Interlending OSI Network" (ISO 10161 and 10162) to connect their effort to similar projects in The Netherlands and the UK.

continued on next page

Europe, at least, discovers the users *(continued)*

In document delivery, she mentioned the now-ended "FOUDRE" project, which used digital scanning and attempted to capture and store text, as it was scanned, for future digital use: this ran into both money and copyright problems. A newer "EDIL/Electronic Document Interchange for Libraries" project, with the UK, Germany, the Netherlands, and Portugal is proceeding, although there still are copyright problems, she said.

Users

Jill Foster, one of the Program Committee members, emphasized the PC background of users, in her presentations on "User Support." She mentioned the large and expanding work group on User Support, "RARE ISUS WG," which now draws from many different international groups, and itself supports user-support work groups in fields as diverse as cetacean studies, developmental psychiatry, diabetes studies, and marine technology. An excellent report edited by Foster, which arrived belatedly after wrestling its way through Italian customs, summarizes European efforts in the user support area ("User Support and Information Services in the RARE Community—A Status Report," RARE Technical Report 1, RARE Working Group 3, Subgroup USIS, 1st edition, March, 1992). Taking a phrase from her countryman, Lorcan Dempsey, Jill reminded the conference that the networks now, "present users with a flea market, when what's needed is a department store": user support badly needs such network organizing, she said.

User support at Cornell

Carole Lambert, from Cornell University, described the hard-nosed managerial analysis to which they subjected their local version of the "computer-center-versus-library" competition in information provision which plagues every campus. "We hit the wall," she said, "with a service that wouldn't scale":

- Systems consulting services were one-on-one,
- Classroom training focused on skills rather than on use and resources,
- Documentation was labor-intensive, with limited distribution,
- Accounting went to and not always through a central bottleneck.

Their new model, she said, presents a "scalable method of delivery": they decided to:

- Leverage the technology—use their own computer and network technology to develop and disseminate user tools,
- Leverage the human resources—they build campus coalitions, shared solutions, use e-mail and other techniques to sustain campus contacts, and "eliminate redundancies" (dangerous-sounding term, I think, for cutting out duplications).

Most of all though, at Cornell, they are trying to change the attitudes and expectations of the users: "we want to make independence easier than dependence," said Lambert, "we teach the users problem diagnosis and resolution along with traditional user skills...we will be there, but we want them to rely on themselves more than they rely on us."

Cost—an idea whose time is about to arrive

Thomas Johannsen, originally of Dresden and now of just-north-of-Tokyo, made a fascinating presentation of *SoftPages,* his "distributed database for fileserver contents" (e.g., *Archie, WAIS*), which has a built-in module for computing usage "cost," in terms of "economic distance"—using speed, tariff, traffic and priority parameters.

Johannsen's presentation struck a chord in the conference: everyone is getting a new awareness of usage costs, as the "academic test-bed" history of the networks recedes and the "commercial" age dawns, and you could see many minds in the audience quickly considering the logistics of building in similar "costing" modules to other tools, following Johannsen's suggestion. (The NREN legislation in the US calls for precisely this sort of new approach: "The Network shall...have accounting mechanisms which allow users or groups of users to be charged for their usage of copyrighted materials..." *High Performance Computing Act* of 1991. t.1,s.102,c,6.)

Comprehensive approach

Willem Scholten presented Project Janus, Columbia University law library's effort to:

- Avoid microforms,
- Bring the library to the user (Manhattan presents critical space problems),
- Adapt to changing patterns of information distribution.

The project involves participation by Thinking Machines Corp., the university's main and health sciences libraries, the law library, and the United Nations library human rights collection. One critical goal was preservation of the law library's unique and rapidly-deteriorating collections of Nuremberg (375,000 double-sided pages) and Rosenberg (250,000 double-sided pages) trial documents. Their solution uses a special *XWAIS,* a highly-customized version of the publicly-available WAIS tool, digitization with OCR, optical and magnetic tape, and Z39.50 and ISO's SR/1, two "Sun Sparc workstation networks," a "Xerox Docutech 7000 scanner and OCR system," and a "CM2-32K Thinking Machines parallel processing super-computer": all the latest stuff.

Many hands have been in on the project: the law school publishes 13 legal periodicals, for example, and the goal of getting such publishing costs back in-house is being approached through SGML and electronic publishing on the system. The reference desk is interested in information which has time value and takes too long to get into print: the system loaded the *North American Free Trade Agreement* recently and at last count was getting 200–250 "hits" per day on that resource, and similar figures have been achieved for on-line versions of the Maastricht Treaty and the papers of the Rio Conference on the Environment. One other library dream, of loading fulltext direct from commercial publishers, also at least is under discussion with Simon & Schuster: user licenses for the library, based on a flat fee with royalties for downloading.

Closing Keynote

John Quarterman began his conference-closing keynote address with the warning that he wouldn't make predictions—"my crystal ball's kinda cloudy," he said—and then proceeded to make them. He has put together a wonderfully-interesting series of maps, all using data taken from various domain-name registries and servers, showing where all the network use is taking place in the world (surprising activity patterns in Iceland, Australia, Moscow, Hong Kong), and suggesting a continuing rate of usage growth so phenomenal as to be catastrophic for both the networks and librarians. It seems still that only Quarterman, despite his good influence exerted since the 1988 publication of his book, *The Matrix,* has the breadth of vision, and the patience, to look at all the world's information networks—Internet, EARN, BITNET, etc.—as a whole.

continued on next page

Europe, at least, discovers the users *(continued)*

Conclusions

• *The impending invasion of the commercial market:* The real problem, lurking behind most of the conference talk, is what to do about the impending invasion of the commercial market. The commercial publishers are poised to plunge into the little world of academic networking, we heard again and again. Quarterman showed us a fascinating map, on which the portion of world network use devoted to "purely academic" activity—which represented *all* network use a short time ago—now is small and is shrinking rapidly: "academic use" will become an insignificant part of networking as a whole, shortly, he asserted.

• *"The academic model will not scale":* The problem, then, acknowledged again and again by US and non-US attendees, is that "the academic model will not scale": as network use grows, the tools and structures and carefully-developed "standards"—think of MARC, SGML, FTP, *telnet,* opaque user interfaces, even ASCII—will not satisfy a non-academic, international, user public: a despairing conclusion—one which left several librarian-users in the audience feeling a little abandoned. "Information overload" then, inevitably, was debated: several people felt that a bad network situation in this respect is about to get much, much, worse.

• *"The academic model had BETTER scale":* Some braver souls, though, insisted that private industry will need some standards as well: if not necessarily for sharing information easily as an altruistic goal, as the academic world wants, then at least for ensuring the compatibility of its own hardware, software, and services with a particular marketing structure: IBM products and services talking to each other, Siemens doing the same, all the components of a local- or wide-area network—serving fulltext newspapers to northern California, Shakespeare to the entire Ivy League, or Montaigne (or Simenon) to Touraine—able to communicate among themselves. Private industry will have to start somewhere in all this, and that beginning may well be made with at least some of the elaborate tools and standards which have been assembled by the careful academic community today. Such, at least, is the hope.

• *The Atlantic is a very wide pond:* It was very interesting for this American to note the fundamental difference between the US and the European approaches on the standards point. Much good work on standards is being done on both sides of the Atlantic. But the intense preoccupation with standards and consensus-building in general is markedly different in Europe than it is in the US. Great levels of bureaucracy, much tedious negotiation, and great levels of frustration, all are devoted to accomplishing the smallest point of agreement in Europe, ruled constantly by the conviction that without some sort of "top-down" agreement, no "bottom-up" effort will succeed. Not that bureaucracy and haggling don't take place in the US context; but there seems to be more in Europe, and it's much more intense, and deemed to be much more necessary. Law students everywhere learn that Anglo-American law may be built piecemeal, upon the "Common Law," and upon individual cases, while Continental law is a seamless web of "codes," which are thought to cover all conceivable instances. There is this same longing for "codification" in European networking standards work: piecemeal, such as has characterized the evolution of the Internet, will not do in Europe—they need "top-down" codes and standards, before they can proceed rather than after—a major difference from the US approach.

Program Committee chairman, Dennis Jennings mused about this, pointedly, to the several US attendees and speakers: "You must remember that you are one, gigantic, country, while we are by comparison a very large, but still very disunited, collection of very small countries." It is interesting to consider whether the US or the European "consensus-model" will more readily "scale up" to the rapidly-evolving world information Matrix.

Information overload

Network Services '92, then, came to the conclusion that it may well become impossible to service the networks during the next few years: too many, too much, understood and aided by too few. The glass which appears half empty, however, also is half full. There will be many more users and many more things to do. Dennis Jennings also pointed out, however, that the evolution of the telephone was aided by a paradigm shift: fears early in this century that there never would be enough telephone operators were answered by the users becoming the operators themselves. Just so, Jennings insists, a paradigm shift will occur in networked information. The bottlenecks which exist today—of costs and hardware capacities and user training and clumsy interfaces—may be resolved ultimately by similar shifts: "transparent" interfaces, "invisible" technologies, "paperless" libraries, "hypertext" organization and access—it's hard to tell what from here, but something.

A role for librarians

One final optimistic note sounded by the conference left the librarians in the audience feeling smug. Already, no one can *find* anything on the "nets," and it seems that this problem is not going away: it seems, in fact, that the entry of the commercial market is about to make the "navigating" problem much, much, worse. Navigating through information resources is what librarians do: it's what we've done for centuries. It is nice to feel needed: it's reassuring to discover how badly we're going to be needed by the information network users in Europe and elsewhere during the next few years.

References

[1] Kahle, Brewster, "An Information System for Corporate Users: Wide Area Information Servers," *ConneXions,* Volume 5, No. 11, November, 1991.

[2] Deutsch, P. & Emtage, A., "The *archie* System: An Internet Electronic Directory Service," *ConneXions,* Volume 6, No. 2, February 1992.

[3] Berners-Lee, T., "A Summary of the WorldWideWeb System," *ConneXions,* Volume 6, No. 7, July 1992.

[4] "Réseaux Associés pour la Recherche Européenne (RARE)," *ConneXions,* Volume 6, No. 1, January 1992.

JACK KESSLER has academic degrees in philosophy, law, and library and information studies, and has pursued these and other subjects at Yale, Oxford, and the University of California. He spent fifteen years in the handicraft importing business, until he found the glamor of international travel to be at odds with the joys of married life and the raising of two small boys. His love affair with books and love/hate relationship with the computer are long-standing. He fought the automation battles of the 70s and 80s, most often siding with the Luddites against the machines but then reluctantly giving in. He's still suspicious. He has been a networked information consultant, a lecturer, an indexer, and a researcher in the law school library at UC Berkeley, and recently he was enrolled at Berkeley, studying the US Internet and the French Minitel. During 1992–3 he has lived in Lyon, France, writing articles and a book. He'll return to his San Francisco home in the summer of 1993. He is a member of the American Society for Information Science, the American Library Association, and the California Library Association, and he spends too much time online on the PACS-L and EXLIBRIS e-conferences and on the WELL. His ambition in life is never to take another airplane trip. E-mail: `kessler@well.sf.ca.us`

Call for Papers

The *Internet Society Symposium on Network and Distributed System Security* will be held 3–4 February, 1994 at the Catamaran Hotel in San Diego, California.

Topics

The symposium will bring together people who are building software and hardware to provide network or distributed system security services. The symposium is intended for those interested in practical aspects of network and distributed system security, rather than in theory. Proceedings will be published by the Internet Society. Topics for the symposium include, but are not limited to, the following:

- Design and implementation of services: access control, authentication, availability, confidentiality, integrity, and non-repudiation—including criteria for placing services at particular protocol layers.
- Design and implementation of security mechanisms and support services: encipherment and key management systems, authorization and audit systems, and intrusion detection systems.
- Requirements and architectures for distributed applications and network functions—message handling, file transport, remote file access, directories, time synchronization, interactive sessions, remote data base management and access, routing, voice and video multicast and conferencing, news groups, network management, boot services, mobile computing, and remote I/O.
- Special issues and problems in security architecture, such as very large systems like the international Internet, and high-speed systems like the gigabit testbeds now being built.
- Interplay between security goals and other goals: efficiency, reliability, interoperability, resource sharing, and low cost.

Submissions

The program committee seeks both original technical papers and proposals for panel discussions on technical and other topics of general interest. Technical papers should be 10–20 pages in length. Panels should include three or four speakers. A panel proposal must name the panel chair, include a one-page topic introduction authored by the chair, and also include one-page position summaries authored by each speaker. Both the technical papers and the panel papers will appear in the proceedings. Submissions should be made via e-mail to:

`1994symposium@smiley.mitre.org.`

Formats

Submissions may be in either of two formats: ASCII or *PostScript*. If the committee is unable to read (print) a *PostScript* submission, it will be returned and ASCII requested. If electronic submission is absolutely impossible, submissions should be sent via postal mail to:

Robert W. Shirey, Mail Stop Z202
The MITRE Corporation
McLean, Virginia 22102–3481
USA
`Shirey@MITRE.org`

All submissions must include both an Internet electronic mail address and a postal address.

Important dates

Submissions due:	August 16, 1993
Notification of acceptance:	October 15, 1993
Camera-ready copy due:	November 15, 1993

Call for Papers

The *Conference on High-Speed Networking and Multimedia Applications,* part of IS&T/SPIE Symposium on Electronic Imaging: Science & Technology, will be held in the San Jose Convention Center, San Jose, California, February 6–10, 1994.

The emergence of reasonably priced high-speed networks that can link ever smaller and more powerful computers creates new opportunities and challenges for users and application developers. Along with these developments, new technologies such as multimedia, visualization, and virtual reality have emerged. The growth of these technologies is proof that the pace of expansion in computer engineering continues to accelerate into the 90s.

This conference intends to bring together researchers and developers who are exploring or building high speed networks and applications enabled by high speed networks. The conference will serve as a forum for discussing advanced networking technologies, novel application ideas, insights in application system modeling, integration and network requirements, and experiences from prototypes, field trials, or production systems. Equal emphases will be placed on network technologies, application systems, and their mutual dependencies.

Topics

Papers are solicited in all areas of high speed networks and multimedia applications, including but not limited to:

- High speed networks
- Collaboration over networks
- Human interface for multimedia applications
- Network based medical imaging
- Multimedia scientific applications
- Distributed multimedia applications
- Games and entertainment
- Distributed virtual worlds
- Security in distributed applications

Submissions

Send a 500 word abstract of your paper to:

S&T/SPIE EI94
1000 20th St.,
Bellingham, WA 98225.
Phone: +1 206-676 3290
Fax: +1 206-647 1445
E-mail: abstracts@mom.spie.org

Submissions should include: Title of Abstract/Paper, authors' full names and affiliations as they will appear in the program, complete mailing addresses, telephone and telefax numbers, and e-mail addresses, plus a brief biography (50–100 words, principal author only).

Important dates

Abstracts due:	July 19, 1993
Acceptance notification:	October 21, 1993
Camera-ready abstract due:	November 15, 1993
Manuscript due:	January 10, 1994

Letter from The President and Vice President in Announcement of White House E-Mail Access

June 1, 1993

Dear Friends:

Part of our commitment to change is to keep the White House in step with today's changing technology. As we move ahead into the twenty-first century, we must have a government that can show the way and lead by example. Today, we are pleased to announce that for the first time in history, the White House will be connected to you via electronic mail. Electronic mail will bring the Presidency and this Administration closer and make it more accessible to the people.

Improved access

The White House will be connected to the Internet as well as several on-line commercial vendors, thus making us more accessible and more in touch with people across this country. We will not be alone in this venture. Congress is also getting involved, and an exciting announcement regarding electronic mail is expected to come from the House of Representatives tomorrow.

Various government agencies also will be taking part in the near future. *Americans Communicating Electronically* is a project developed by several government agencies to coordinate and improve access to the nation's educational and information assets and resources. This will be done through interactive communications such as electronic mail, and brought to people who do not have ready access to a computer.

Limitations

However, we must be realistic about the limitations and expectations of the White House electronic mail system. This experiment is the first-ever e-mail project done on such a large scale. As we work to reinvent government and streamline our processes, the e-mail project can help to put us on the leading edge of progress.

Initially, your e-mail message will be read and receipt immediately acknowledged. A careful count will be taken on the number received as well as the subject of each message. However, the White House is not yet capable of sending back a tailored response via electronic mail. We are hoping this will happen by the end of the year.

A number of response-based programs which allow technology to help us read your message more effectively, and, eventually respond to you electronically in a timely fashion will be tried out as well. These programs will change periodically as we experiment with the best way to handle electronic mail from the public. Since this has never been tried before, it is important to allow for some flexibility in the system in these first stages. We welcome your suggestions.

Historic

This is an historic moment in the White House and we look forward to your participation and enthusiasm for this milestone event. We eagerly anticipate the day when electronic mail from the public is an integral and normal part of the White House communications system.

President Clinton	`president@whitehouse.gov`
Vice President Gore	`vice.president@whitehouse.gov`

House Announces Public Electronic Mail Service

(June 3, 1993) Chairman Charlie Rose and Ranking Minority Member Bill Thomas of the Committee on House Administration announced today the pilot program of the *Constituent Electronic Mail System.* This ground-breaking new service will allow citizens to communicate directly with their Member of Congress by electronic mail. The House of Representatives has established an electronic gateway to the Internet, the vast computer network that is used currently by over twelve million people worldwide. Participating Members of the House have been assigned public mailboxes which may be accessed by their constituents from their home computers. In addition, many libraries, schools and other public institutions now provide, or soon will provide, public access to the Internet.

Participants

The Members of the House of Representatives who have agreed to participate in this pilot program are: Rep. Jay Dickey (AR-07), Rep. Sam Gejdenson (CT-02), Rep. Newt Gingrich (GA-06), Rep. George Miller (CA-07), Rep. Charlie Rose (NC-07), Rep. Fortney (Pete) Stark (CA-l3), and Rep. Melvin Watt (NC-12). These Members will be making announcements in their congressional districts within the next few weeks to make their constituents aware of the new service.

The *Constituent Electronic Mail System* represents a significant effort by the House of Representatives to expand communication with constituents. With the tremendous growth of electronic mail over the past several years, and the increasingly inter-connected nature of computer networks, the new service is a natural addition to the current methods of communication available to constituents. At the present time, House Members involved in the pilot program will largely respond to electronic mail messages from their constituents by postal mail, to ensure confidentiality.

Constituents of House Members participating in the pilot program who wish to communicate with those Members will be asked to send a letter or postcard stating their interest to the Member's office. The request will include the constituent's Internet "address," as well as that constituent's name and postal address. This process will allow Members to identify an electronic mail user as his or her constituent.

The pilot e-mail program will continue until sufficient feedback from participating offices has been collected to allow improvements and modifications to the system. When House Information Systems and the Committee on House Administration are satisfied that the system is sufficiently error-free, other Members of the House will be allowed to add this new service as technical, budgetary and staffing concerns allow.

More information

For more information, Internet users are encouraged to contact the House of Representative's new on-line information service. Please send a request for information to `congress@hr.house.gov` or call Lance Koonce at +1 202-225-7922.

[Ed.: We don't normally reproduce press releases verbatim in this publication, but the importance of these two announcements seems to warrant a departure from normal operating procedures.]

Interop Announces NetWorld+INTEROP 94 Berlin

Interop Company has recently announced *NetWorld®+INTEROP® 94 Berlin,* a combination conference and exhibition focused on the technical and educational aspects of integrating all levels of communications and computing. The new event will be held from June 6–10, 1994 at the AMK Fairgrounds and adjacent ICC Convention Center in Berlin, Germany. Its goal is to provide both management and technical staff in European-based businesses and government offices with the opportunity to learn about the latest networking technologies and products that will help strengthen their organizations for today's rapidly changing global economic and political conditions.

This announcement of a Berlin venue is another step in plans first unveiled by Interop Company and Novell, Inc. on December 18, 1992. At that time, the two companies announced the creation of an event called *NetWorld+INTEROP 94,* a single forum for addressing the industry's growing need worldwide to understand interoperability issues at all levels of networking and computing—from the desktop to the data center. In January 1993, Interop announced that NetWorld+ INTEROP 94 will be held in the United States May 2–6, 1994 in Las Vegas, Nevada and in Atlanta, Georgia from September 12–14, 1994.

A new formula for European audiences

NetWorld+INTEROP 94 Berlin is Interop Company's second event outside the United States. Like the first, INTEROP® Europe 93, which will be held in Paris, France, October 25–29, 1993, the conference portion of the Berlin event will be managed by a program committee of local industry experts with in-depth knowledge of networking standards, protocols and other technology issues relevant to European audiences. In addition to the exhibition and general conference, the five-day Berlin event will offer attendees a choice of special conferences, in-depth tutorials led by European and American academic experts, the INTEROPnet™ network and Solutions Showcase™ demonstrations highlighting the practical applications of new technologies.

"In Europe, trade shows and academic conferences are usually held separately and target different audiences," said Daniel C. Lynch, president of Interop Company. "That situation is changing, however. As in the United States, European business and government leaders realize that interoperability is integral to their computing and networking purchase decisions. They seek a forum where they can study the complex issues of internetworking, learn about the available solutions and see how different vendors' products work together. In the past year, we have seen a growing number of European attendees at Interop Company events in the U.S., asking when we plan to bring the 'Interop experience' closer to their home territory. First with Paris and now with Berlin, we're responding to those requests."

About Interop Company

Interop Company is an educational services and conference and exhibition company specializing in computers and communications. Interop is dedicated to bringing the latest ideas from researchers, analysts and vendors to the user community through conferences, seminars, publications and educational services. Founded in 1985, Interop Company is the sponsor of conferences and expositions worldwide, including NetWorld+INTEROP 94 and INTEROP, the premier forum for addressing the interoperability challenges and solutions found in the real world of enterprise computing.

New Internet Society Trustees Elected

Introduction

The Internet Society (ISOC) completed its first election of new trustees on June 1, 1993 adding six trustees to the existing thirteen. Newly-elected to the ISOC Board of Trustees are:

Scott Bradner: Consultant, Harvard University Office for Information Systems.

Susan Estrada: Consultant, President of the Commercial Internet Exchange (CIX) and a founder of the California Education and Research Federation Network (CERFnet).

David Farber: Professor of Computer and Information Science and Director, Distributed Systems Laboratory, University of Pennsylvania.

Haruhisa Ishida: Professor, University of Tokyo and Head of Research and Development, Computer Center.

Jean Armour Polly: Manager of Network Development and User Training, NYSERNet.

Jonathan Postel: Associate Director for Networking of the High Performance Computing and Communication Division of the Information Sciences Institute of the University of Southern California.

They join the current trustees in guiding the Internet Society's growth and direction.

About the society

The Internet Society is a scientific and educational professional membership organization devoted to promoting the evolution and use of the Internet and its technology on a world-wide basis. Founded in 1992, ISOC provides the organizational framework for the development of Internet Standards through the Internet Architecture Board and the Internet Engineering Task Force. The Internet is a system of nearly 12,000 interlinked computer networks containing over 1.7 million host computers in all sectors of the economy (government, academia, education and industry) in over 100 countries. The system has been doubling in size annually for the last several years.

"The new trustees are representative of the vigorous, creative and rapidly growing Internet community. Their involvement in the ISOC Board of Trustees will enhance its perspectives and increase its already considerable energy levels," said Vinton Cerf, President of the Internet Society.

The new trustees will meet face to face with their colleagues on the Board at the 1993 annual conference of the Internet Society, INET '93, August 17–20, in San Francisco.

More information

Background material about the Internet Society, including a list of current trustees can be obtained from the secretariat:

Internet Society
1895 Preston White Drive, Suite 100
Reston, VA 22091
Phone: +1 703-648-9888
Fax: +1 703-620-0913
Internet: `isoc@isoc.org`

Network Reading List available

by Charles Spurgeon, University of Texas at Austin

A new version of the document "Network Reading List: TCP/IP, UNIX, and Ethernet" is now available from the Network Information Center at the University of Texas at Austin. The list may be found on host `ftp.utexas.edu` (128.83.185.16). This is version 4.0 of the reading list, dated June, 1993.

What is it?

The network reading list is an annotated list of books and other resources focusing on three networking technologies that are in wide use: TCP/IP, UNIX, and Ethernet. A mix of resources is presented ranging from introductory information to in-depth technical details.

Version 4.0 of the list has been completely rewritten and updated, and now includes over 80 items. The list is weighted towards resources that cover the territory well, and that deal with real-world problems found on growing networks. A copy of the table of contents is included below.

How do I get a copy?

You can retrieve a copy of the list in either *PostScript* format or as a plain ASCII text file. The *PostScript* format is recommended. Copies of the list may be retrieved using anonymous FTP or a mail-based archive server program.

FTP

The hostname for anonymous FTP is `ftp.utexas.edu`, and the files are in the `pub/netinfo/docs` and `pub/netinfo/ps` directories as `net-read.txt` and `net-read.ps` respectively.

E-mail

The e-mail address of the archive server program is `archive-server@ftp.utexas.edu`. You can retrieve copies of the list by sending the archive server program a command line in the body of an electronic mail message. The command line:

```
send ps net-read.ps
```

will cause the archive server program to send you a copy of the *PostScript* file, while the command line:

```
send docs net-read.txt
```

will retrieve the ASCII text file.

The command line should be placed in the body of an e-mail message sent to the address `archive-server@ftp.utexas.edu`. The archive server is just a simple program, so make sure to send the command line exactly as shown here.

Note: Due to slow links and other problems some mailers will not accept files larger than 100KB. The *PostScript* version of the reading list is over 260KB, and the text version is approximately 154K. If your mailer limits the size of files, you can get around this limitation by retrieving the following files:

```
docs: net-read-1.txt, net-read-2.txt

ps: net-read-1.ps, net-read-2.ps, net-read-3.ps
```

All of these files are smaller than 100KB. You will have to use an editor to restore the *PostScript* files to one large file in order to get the correct output on a *PostScript* printer.

Table of contents

Charles Spurgeon
Campus Networking Services
University of Texas at Austin
Room COM1
Austin, TX 78712
USA
E-mail: c.spurgeon@utexas.edu
Phone: +1 512-471-3241 ext 265

Announcing the Guide to Network Resource Tools

The *Guide to Network Resource Tools* describes many of the key tools in use today among the academic networking community for accessing resources on the net.

Organization

The tools described in this guide have been divided into five functional areas. The first section, "Exploring the Network," covers *Gopher* and *World-Wide Web*. *WAIS* and *ASTRA* are documented in section two, "Searching Databases." The third section, "Finding Network Resources," deals with *Archie, WHOIS* and *NETSERV*. *Trickle* and *BITFTP* are covered in section four on "Getting Files." The final section, "Networked Interest Groups," discusses *Listserv* and *Netnews*.

For each tool, the guide provides a general overview and details on availability, intended audience, basic usage, and examples. Thus, it should complement the work of wg-isus by providing a more detailed picture for those who are interested.

The guide was produced by the EARN staff. It has been printed as an A5 booklet. Comments, corrections and criticisms are encouraged, as EARN plans to produce a second edition of this guide for the *Network Services Conference* (NSC '93) in Warsaw in October.

Getting the guide

The Guide to Network Resource Tools is available electronically from:

```
LISTSERV@EARNCC.BITNET
```

in *PostScript* format.

To get the Guide to Network Resource Tools, send the command:

```
GET NETTOOLS PS
```

Most of the individual sections are available in text format with the command:

```
GET toolname MEMO
```

—David Sitman, EARN

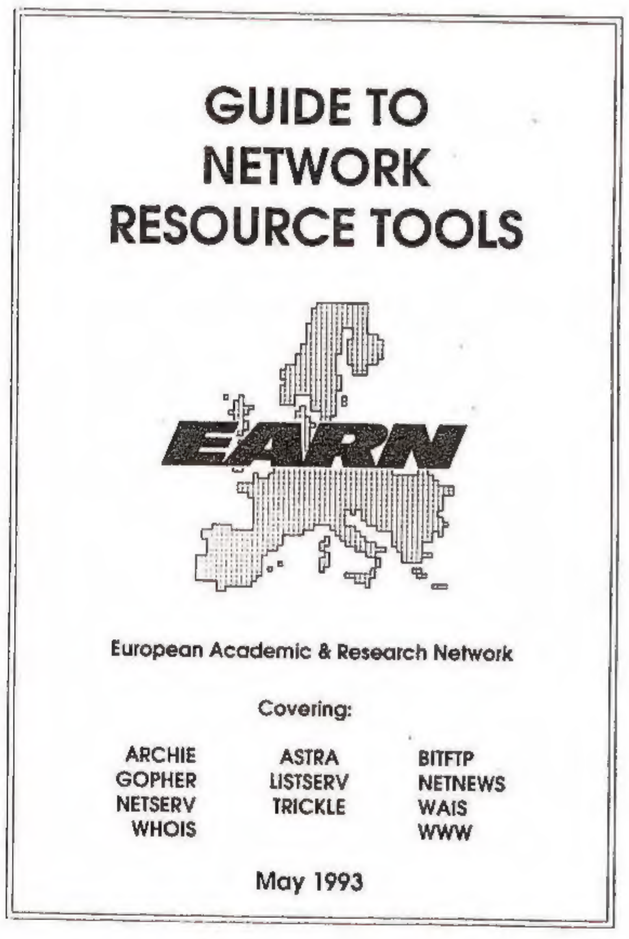

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now over 75 to choose from; ask for our free 1987–1992 index booklet and the 1993 partial index sheet). We want to hear from you. Send your questions, comments or suggestions to:

ConneXions—The Interoperability Report
480 San Antonio Road
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: connexions@interop.com

Don't forget: Next month is INTEROP 93 August in San Francisco! Call 1-800-INTEROP or 1-415-941-3399 for more information or to register.

Journal Review

Useful

I was recently invited to review the first three issues of Volume 1 of the *Journal of High Speed Networks*. (The journal offers a special discount subscription rate of $50 per year to SIGCOMM members.) I accepted the invitation with a bit of reluctance, since I'm not generally in favor of new journals. I think most of us have too many journals to read already and most of those journals aren't worth reading. However, despite this predisposition to dismiss the journal, I actually found it useful and plan to subscribe to it under the SIGCOMM offer (which is described in more detail at the end of this review).

Publishing details

First, a few details about the journal are in order. It is published by IOS Press and its editor-in-chief is Dr. Deepinder Sidhu of the University of Maryland, who hosts an annual gigabit workshop in the DC area, and is this year's program chair for ACM SIGCOMM '93. The editorial board is populated with the usual gigabit suspects such as Acampora, Farber, Gopal, Kleinrock, Maxemchuk, Pehrson, and Spaniol. The journal comes out quarterly, and published about 325 pages (16 papers) in Volume 1.

Sample papers

The first three issues of Volume 1 include a range of papers from theory papers to testbed surveys to architecture papers. Rather than discuss all the papers, this review will briefly present three papers to give the reader some sense of the kinds of papers in the journal.

In Volume 1, Number 1, there is a review of the *LuckyNet* gigabit testbed at AT&T Bell Labs, by Gitlin and London. The paper is interesting both because it surveys a US testbed that is not one of the big six (*Aurora, Blanca, Casa, MAGIC, Nectar,* and *Vistanet*) and because it offers a telephony perspective on the future of gigabit networking, as well as explaining the components of the LuckyNet testbed.

Also in issue Number 1 is a paper on real-time internetwork communication by Ferrari. This paper is an architectural paper, pulling together a lot of interesting work by the *Tenet* group at U.C. Berkeley into a complete framework. While Ferrari and the rest of the group have generally made their architectural goals clear in prior papers, I still found this paper useful to read and a nice summary of their work.

Issue Number 2 contains a very interesting paper by Cruz, extending his work and Parekh's work on performance bounds to cover new service disciplines akin to the *FIFO+* algorithm proposed by Clark, Shenker and Zhang at last year's SIGCOMM conference. While these three papers are among the best in the journal, several of the other papers are comparable in their quality.

My only concern about the papers is that for several of them I'd heard at least parts of the work they present already. (For example, much of the *LuckyNet* work was presented in a series of papers at *GLOBECOM '91*). In each case, the version in the journal was better and more complete that prior versions, but it does leave one with the sense that sometimes one is getting the same material a second time. But, at least so far, the *Journal of High Speed Networks* version seems superior to me.

Subscriptions

If this journal sounds of interest, you can take advantage of the SIGCOMM discount by sending a check or credit card approval for $50 (USD) for a year's subscription to: *Journal of High Speed Networks,* IOS Press, Van Diemenstraat 94, 1013 CN Amsterdam, The Netherlands. Simply note with your request for a subscription that you are a member of SIGCOMM.

—Craig Partridge

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1